Applications and Benchmarking

Analyzing Applications

- The PERC performance tools have been used to collect performance data for a number of applications. This data will drive the modeling and optimization efforts.

Benchmarking

- PERC aims to produce a well-designed, easily configured, targetable set of benchmarks that use the memory hierarchy extensively.
- User communities and projects rather than individual codes are being targeted in order to study the discipline-specific nature of the benchmarks. The focus is on codes that represent the most important tasks which span the variety of codes employed.
- Benchmarking data has been collected for codes from primary application areas on a number of HPC systems.

Interactions with Application Groups

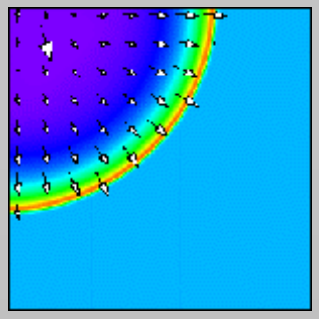

PERC has forged a set of collaborations with an emphasis on application groups in those areas that were stressed in the SciDAC call for proposals.

Collaboration with six SciDAC application groups and two SciDAC ISICs has led to application analysis and benchmarking work in five key areas:

- High Energy and Nuclear Physics – EVH1, Agile-Boltztran, MILC
- Biology and Environmental Research – PCTM, CCM3/MP-2D
- Fusion Energy Sciences – AORSA3D
- Chemical Sciences
- Advanced Scientific Computing - pVarden

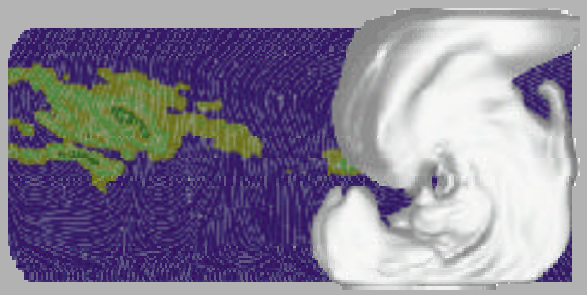

High Energy and Nuclear Physics
EVH1

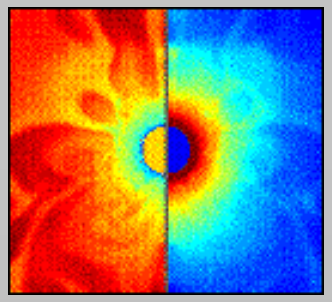

Enhanced Virginia Hydrodynamics #1 represents an important kernel in the "TeraScale Simulations of Neutrino-Driven Supernovae and Their Nucleosynthesis" SciDAC project. The program is currently configured to run a simulation of the Sedov-Taylor blast wave solution in 2D spherical geometry.

Extensive performance studies have been done to collect performance measures on dominant routines, and to study the parallel speedup performance and the communication efficiency of the full code.

Using PAPI and SvPablo, we identified four dominant routines and two candidates for fine-tuning.

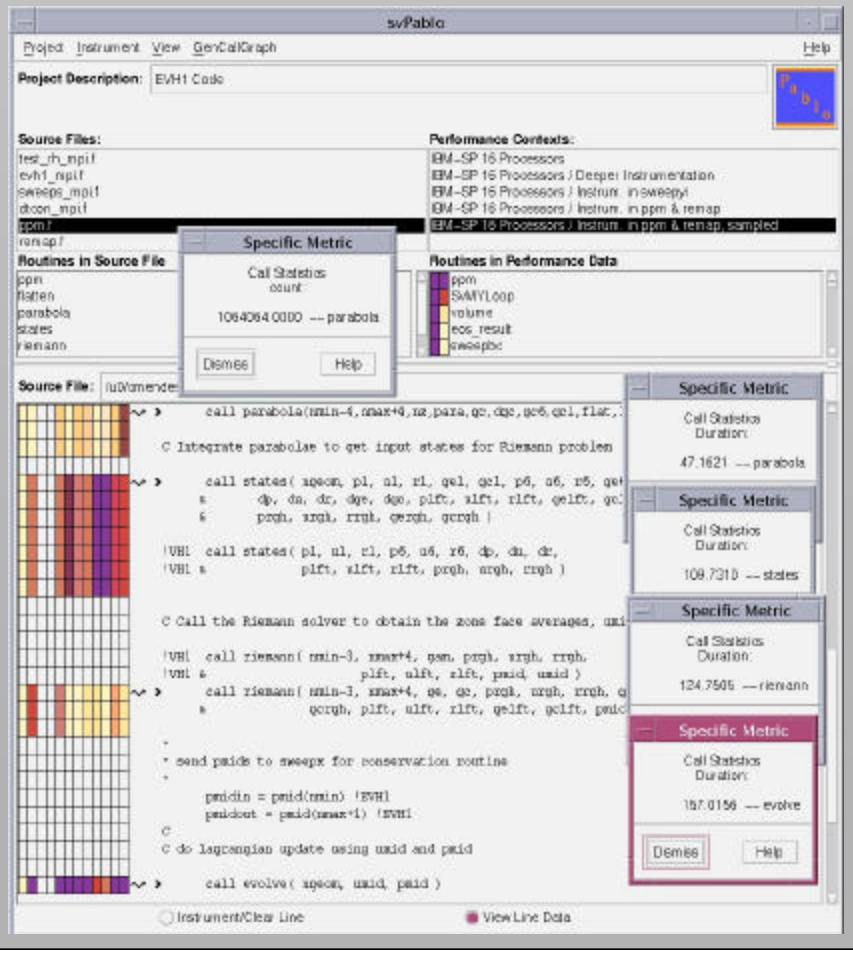

High Energy and Nuclear Physics
EVH1

Aggregate performance measures over all tasks for a .1 simulation-second run. Collected with PAPI on an IBM SP (Nighthawk II / 375MHz).

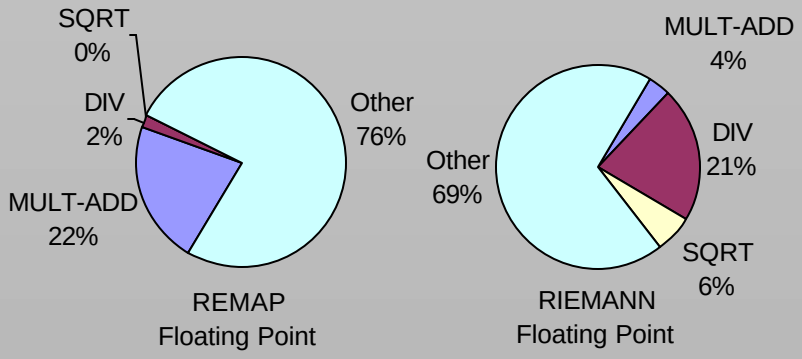

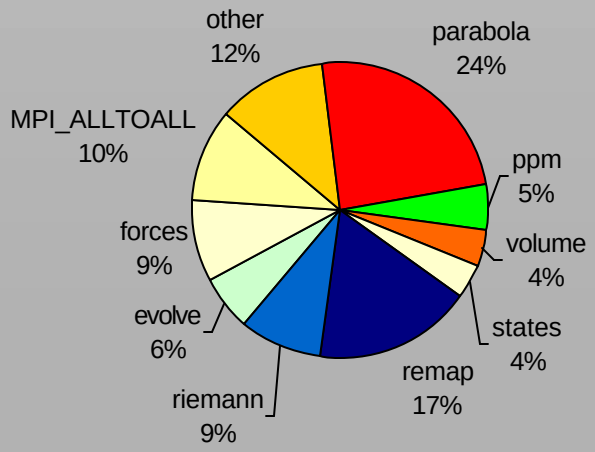

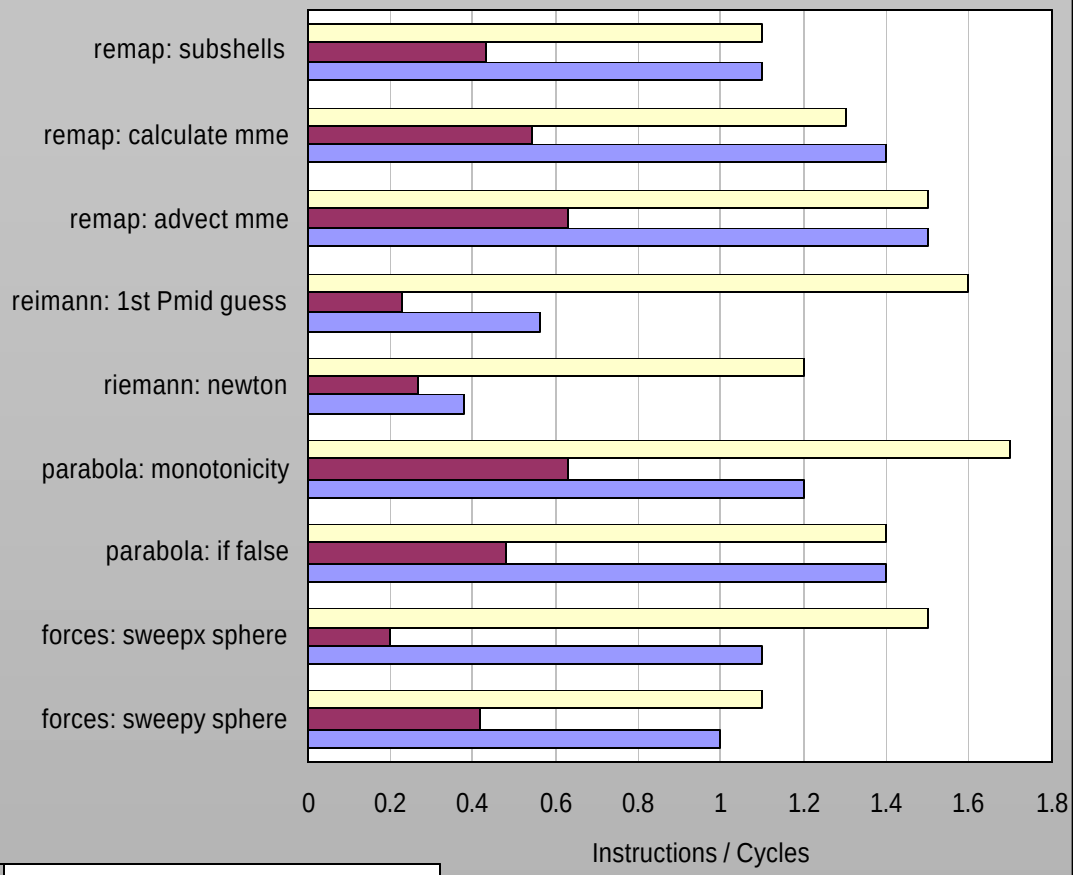

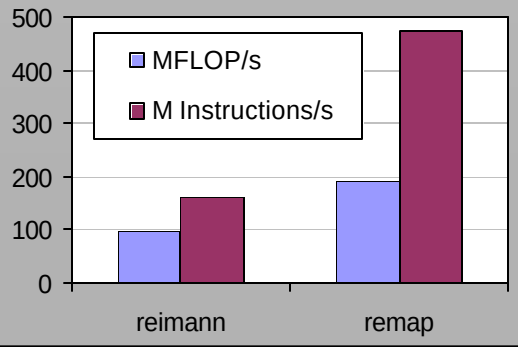

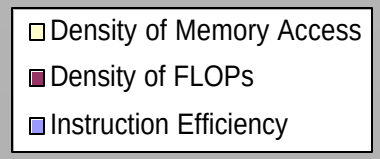

High Energy and Nuclear Physics
EVH1

It was found that for more than 64 processors, EVH1 is communication-bound. The predominant routine (>50% of execution time) at this scale is MPI_ALLTOALL, which is used in matrix-transpose-like operations.

The current implementation uses a 1D decomposition for the matrix-array; a modeling and analysis study has shown that a 2D decomposition would result in a large improvement. Benchmarking and analysis results are below.

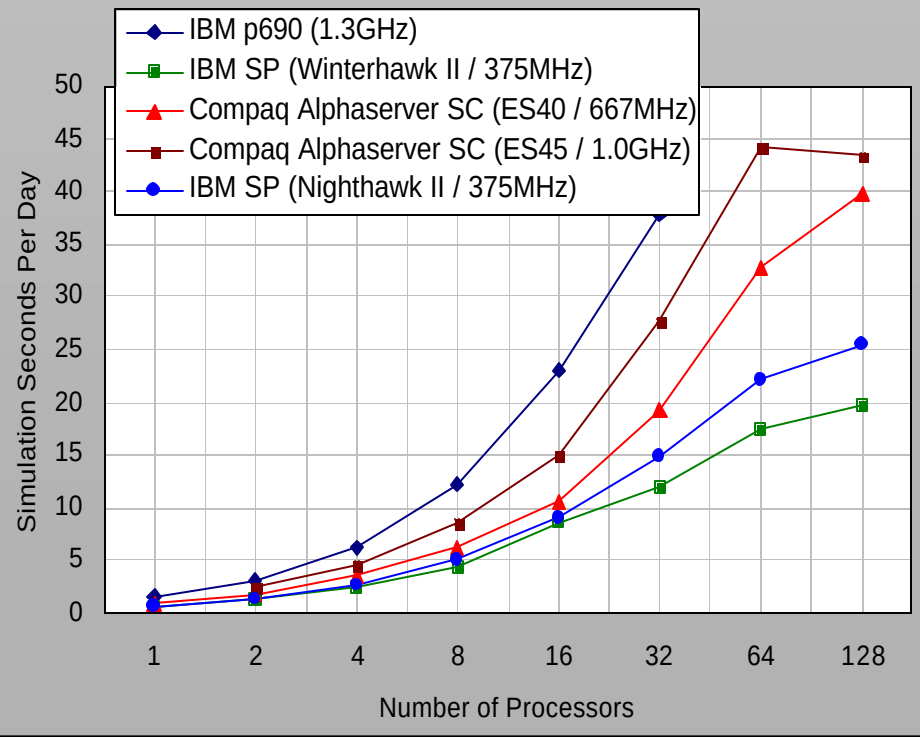

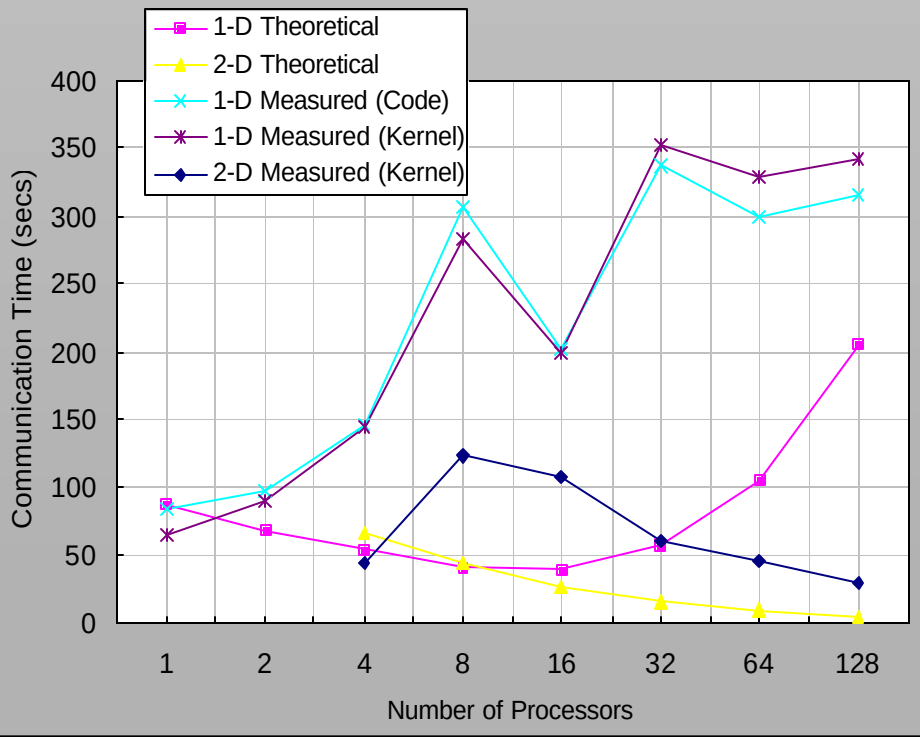

High Energy and Nuclear Physics
AGILE-BOLTZTRAN

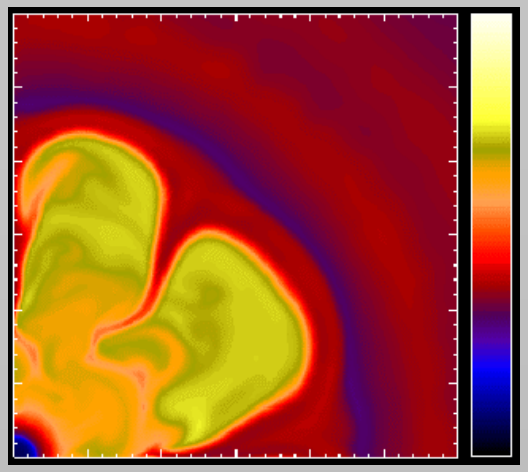

AGILE-BOLTZTRAN is a neutrino radiation hydrodynamics code from the "Terascale Simulations of Neutrino-Driven SuperNovae and Their NucleoSynthesis" SciDAC project. It is used for self-consistent simulations of core-collapse supernovae.

The code incorporates adaptive mesh hydrodynamics and discrete ordinates transport methods in spherical symmetry. Domain decomposition in radius is used, and hydrodynamics is performed redundantly on all nodes. A linear system solve is done via custom ADI preconditioner and various Krylo subspace methods (GMRES, BiCGstab, fixed point iteration).

The results of a preliminary speedup study are shown to the right.

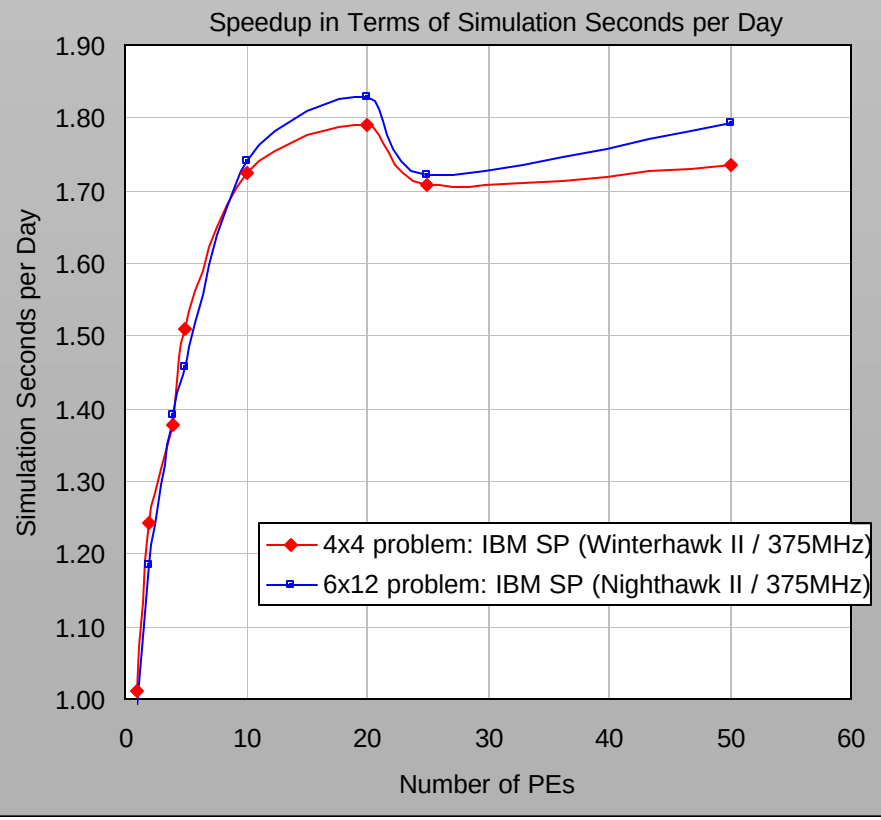

High Energy and Nuclear Physics
MILC

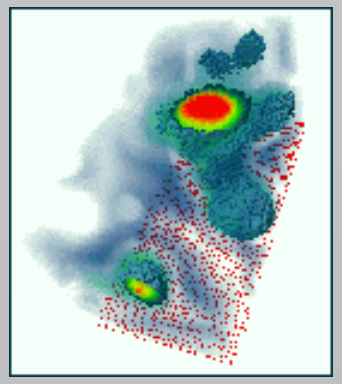

The MILC code, from the National Infrastructure for Lattice Gauge Computing SciDAC project, is a set of codes for doing simulations of four dimensional SU(3) lattice gauge theory on MIMD parallel machines.

The code's computation and communication were studied using SvPablo on executions having between 4 and 32 processors (two 800MHz processors per node) of NCSA's Itanium Linux cluster.

Scalability plots are over $\log_2$(processors):

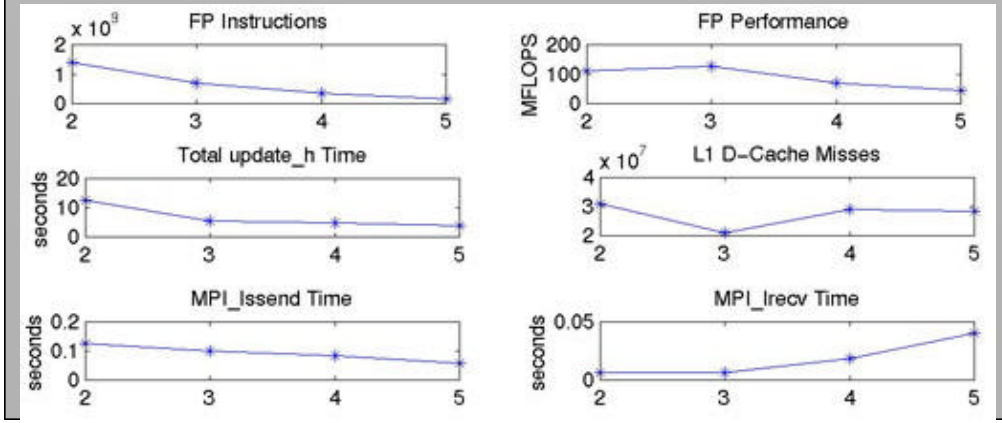

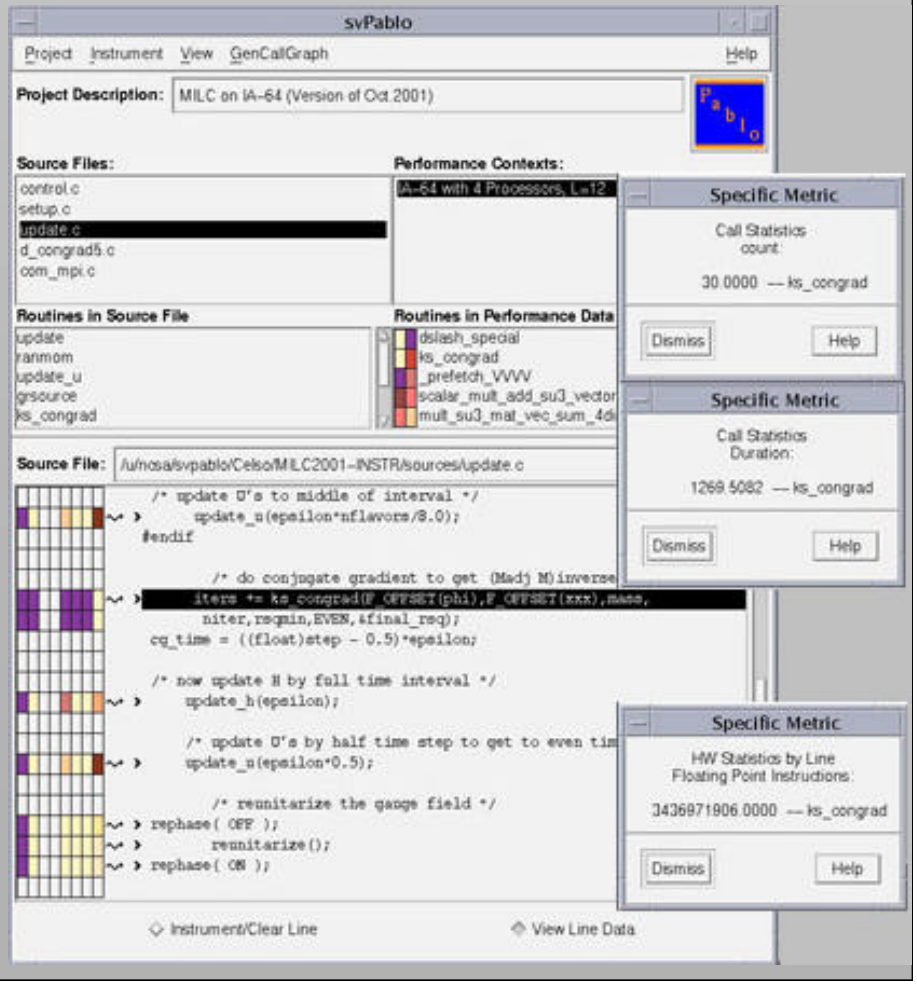

Biology and Environmental Sciences
PCTM

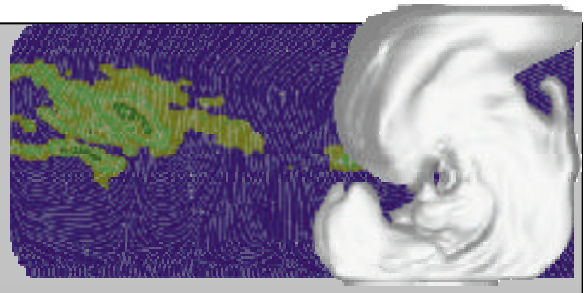

The Parallel Climate Transitional Model (PCTM), from the Climate modeling center, is the next generation of the Parallel Climate Model. It is made up of atmosphere, ocean, land surface, river transport, and sea ice component models, and a coupler to exchange fluxes between the component models.

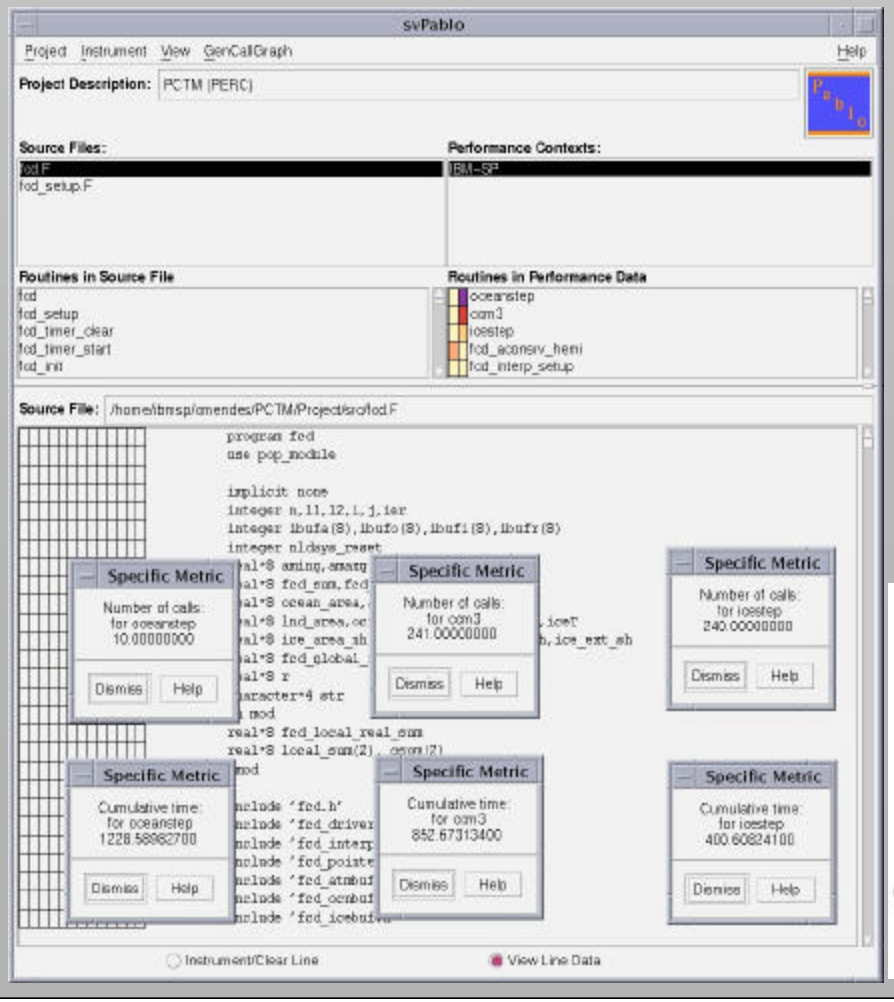

The code has been ported to IBM, Compaq, and Intel platforms. Detailed performance analysis has been done using 4 to 64 processors of an IBM SP using SvPablo.

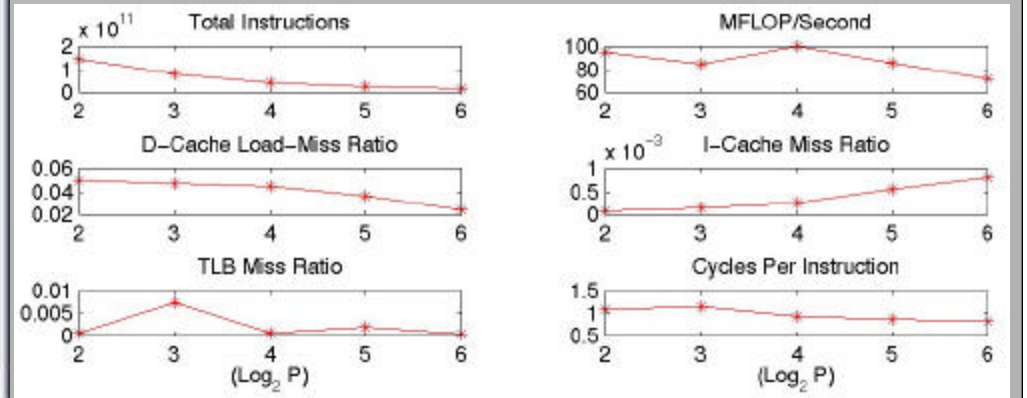

Biology and Environmental Sciences
PCTM

In other studies with PCTM, we examined the impact of different domain decompositions, as well as using fewer MPI processes per SMP node. These studies indicated a strong performance dependence on message-passing performance, where using 64 MPI processes on 32 4-way SMP nodes (leaving 64 processors idle) was 30% faster than when running on 16 SMP nodes.

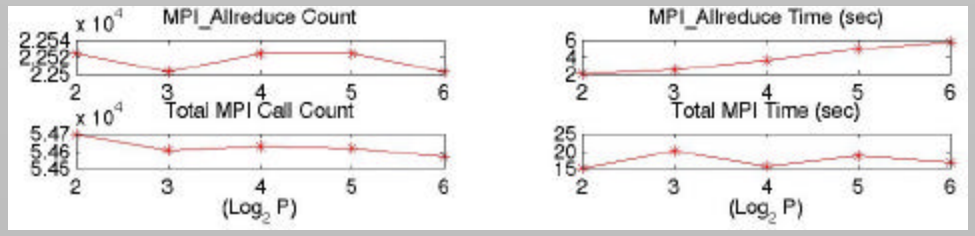

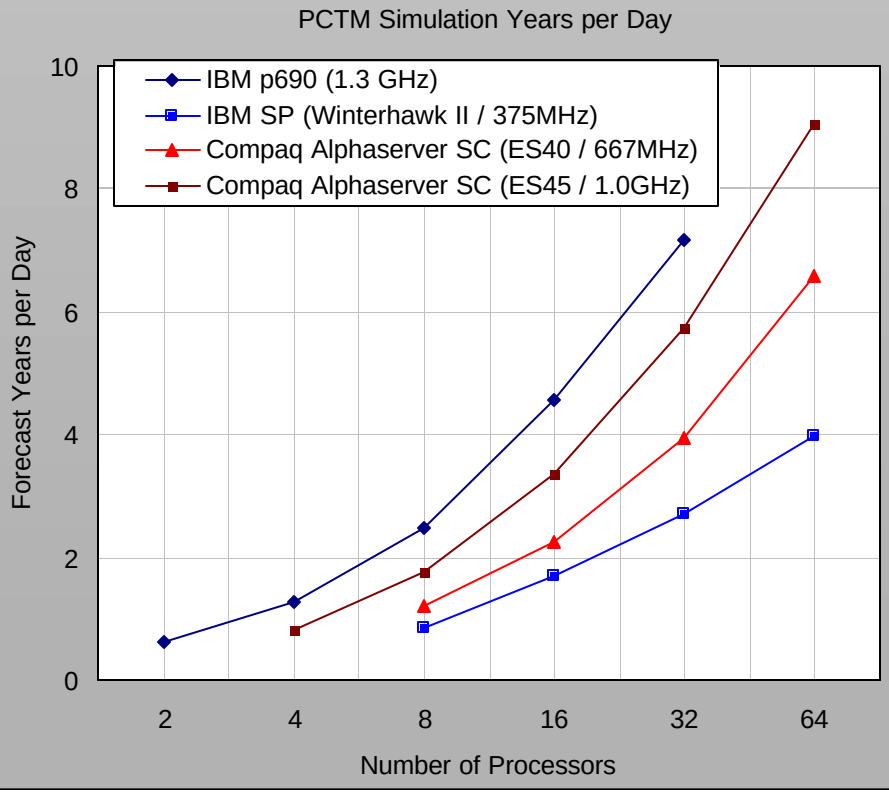

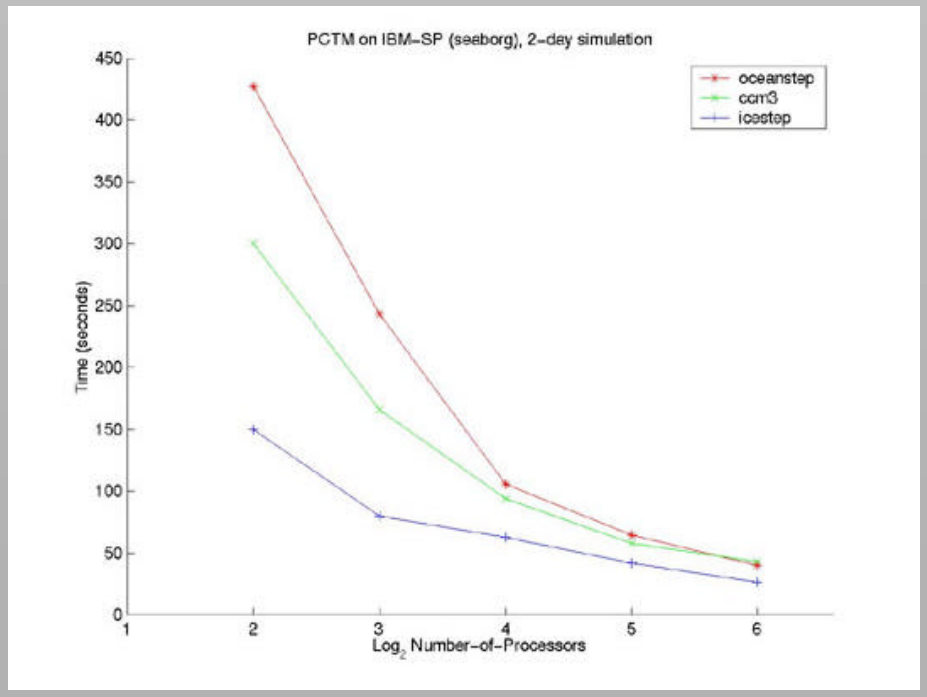

Biology and Environmental Sciences
CCM/MP-2D

CCM/MP-2D is the massively parallel implementation of version 3.6.6 of the Community Climate Model (CCM). It was developed originally to determine how best to parallelize the CCM, and the results from this research are being used in the parallelization of the Community Atmoshperic Model (CAM). CCM/MP-2D is currently used for benchmarking parallel systems. CCM/MP-2D benchmark data for some of our systems is shown below.

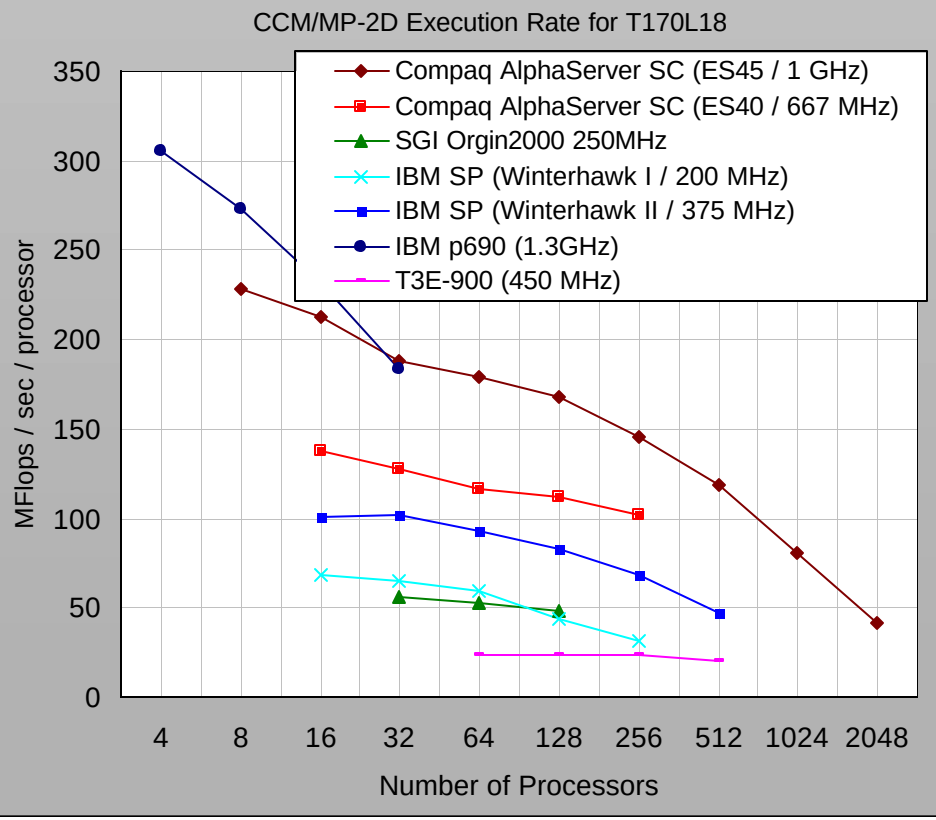

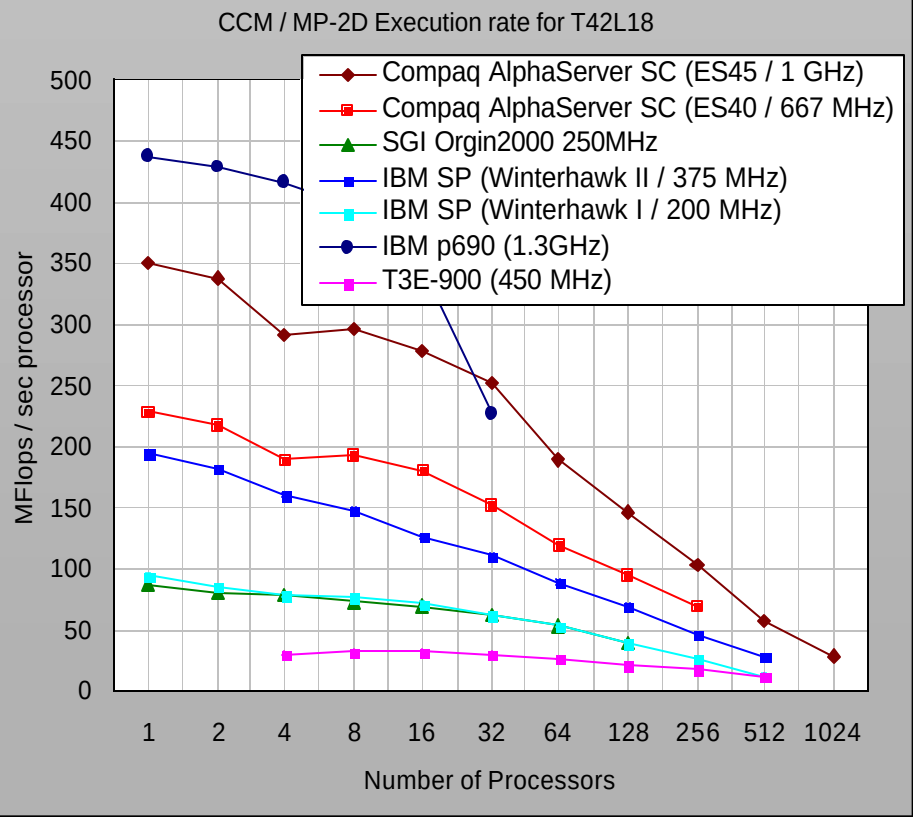

Fusion Sciences
AORSA3D

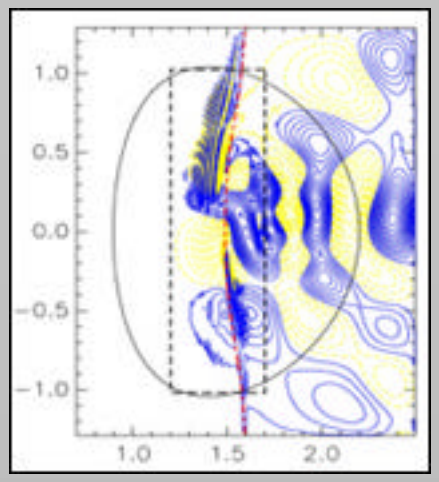

All Orders Spectral Algorithm 3D (AORSA3D) represents an important kernel in the "Numerical Computation of Wave-Plasma Interactions in Multi-dimensional Systems" SciDAC project. The code solves for the wave electric field and heating in a 3-D stellerator plasma heated by radio frequency waves. It is an MPI code that uses SCALAPACK to solve linear systems arising from the spectral discretization.

The code employs ScaLAPACK, which is used to solve a set of dense linear equations. The scaling of the code is determined by problem size and the number of rows and columns used in the block cyclic decomposition (used in ScaLAPACK to distribute the matrix across the processors). It is advantageous to pick block sizes that match as closely as possible to cache memory block sizes used in LAPACK.

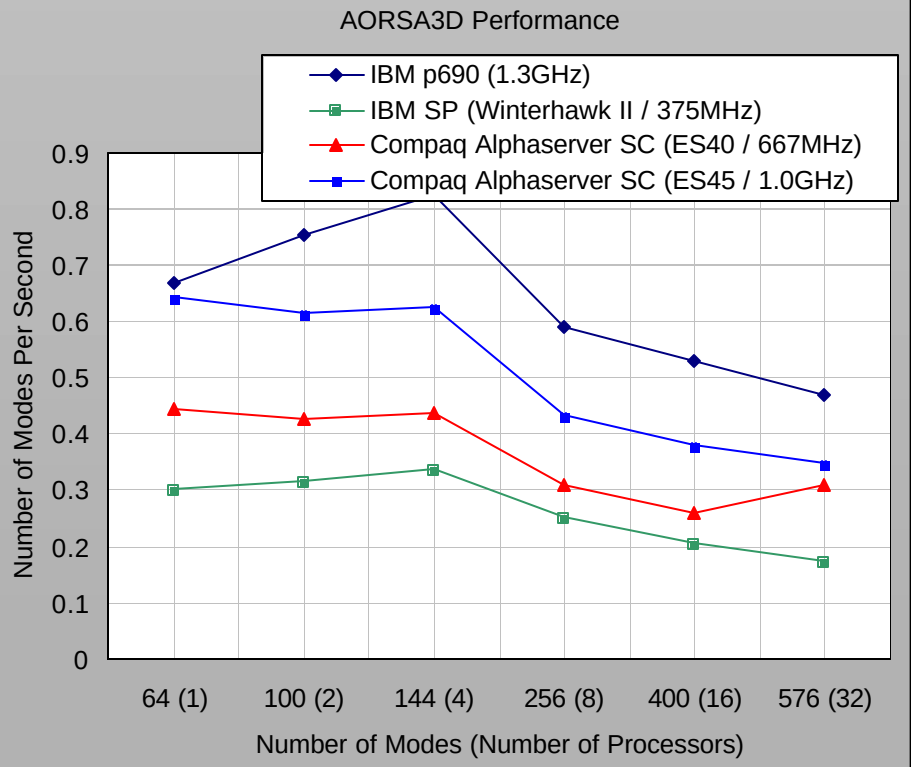

Fusion Sciences
AORSA3D

AORSA3D was ported and benchmarked on IBM and Compaq platforms. A detailed performance analysis has begun using SvPablo and PAPI. The results below are for a 400 Fourier mode run on 16 processors and 1 node of an IBM SP (Nighthawk II / 375MHz).

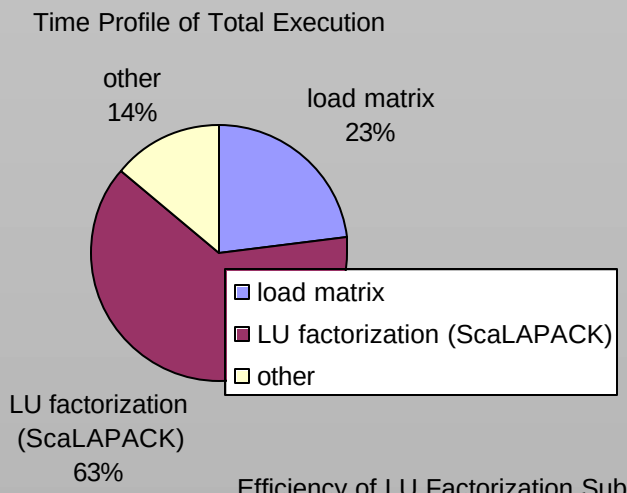

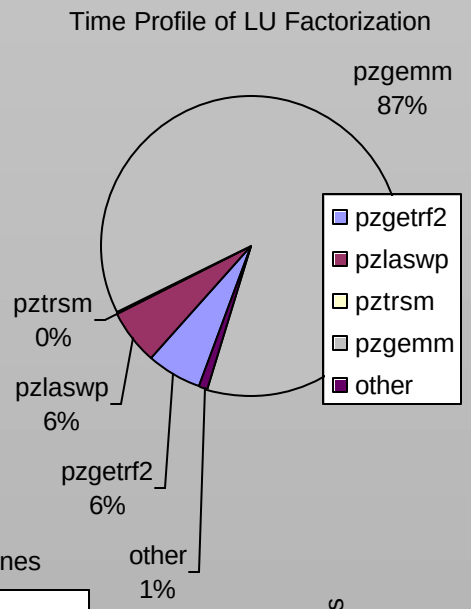

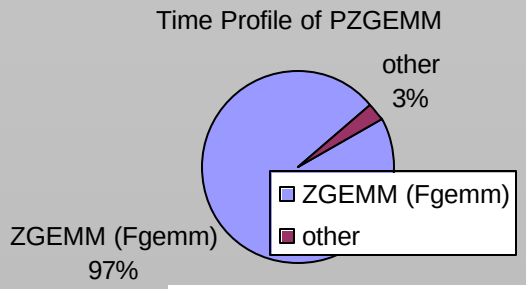

Performance for ZGemm

Denisty of Mem Access	1
Denisty of FLOPs	1.8
MFLOP/s	664
L1 cache hit rate	0.98
L2 cache hit rate	0.96
TLB misses	285034

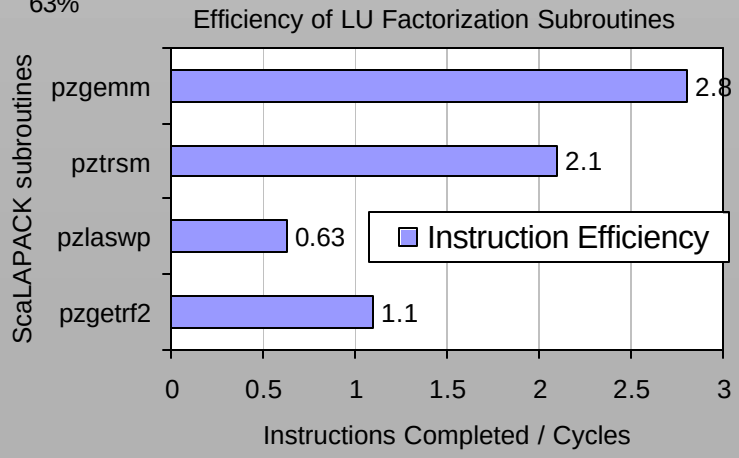

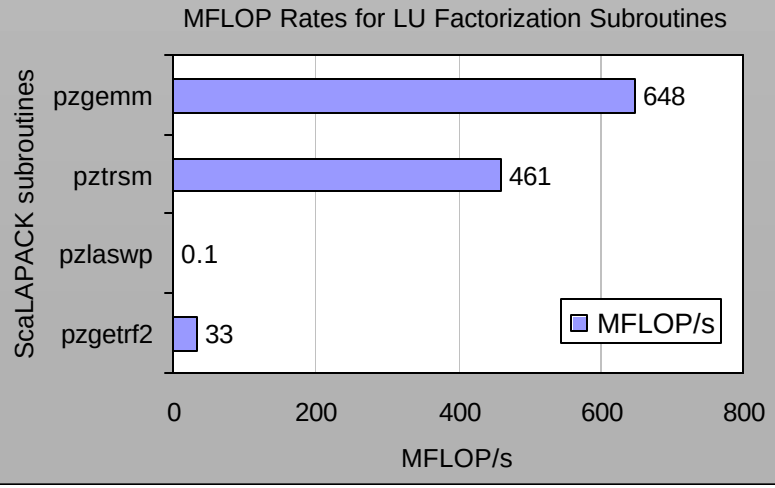

Advanced Scientific Computing
pVarDen

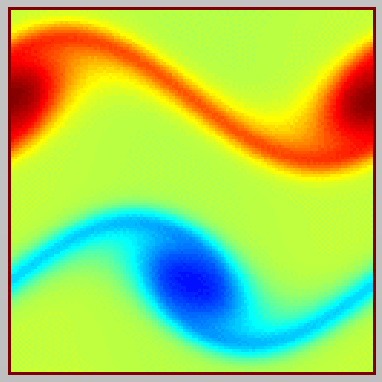

pVarden, or “Parallel VarDen” is an application code from LBNL CCSE for simulating the Variable Density Navier-Stokes equations. It represents an important kernel to the Applied PDEs ISIC. The algorithm used in pVarden is a single grid version of the variable density projection method implemented in CCSE’s IAMR application.

A detailed analysis of pVarden with PAPI has revealed 3-4 dominant routines which vary with the input data set. Some preliminary optimization work has been done to improve the performance of one of these routines.